

12

Describing the Essential Ingredients for an Open, General-Purpose, Shared and Both Large-Scale and Sustainable Experimental Facility (OpenLab)

**Harris Niavis¹, Thanasis Korakis¹, Serge Fdida²
and Loic Baron²**

¹CERTH, University of Thessaly, Greece

²UPMC Sorbonne University, France

12.1 Introduction

OpenLab was an instrumental project, delivering OneLab, the first heterogeneous federation of testbeds open for services on August 2015 and having served hundreds of users since.

OpenLab brought together the essential ingredients for an open, general purpose, and sustainable large-scale shared Future Internet Research and Experimentation (FIRE) Facility, by advancing early prototypes of this Facility. Its main goal was to advance the community by pushing the envelope of a more mature facility, targeting user interface, control and experimentation planes for highly heterogeneous testbeds as well as monitoring and first line support tools. It brought together the most experienced experts and teams in a 3 years project.

The early prototypes, coming from former FIRE initiatives OneLab and Panlab, as well as other valuable sources, included a set of demonstrably successful testbeds: PlanetLab Europe, with its 150 partner/user institutions across Europe; the NITOS and w-iLab.t wireless testbeds; two IMS telco testbeds for exploring merged media distribution; the GSN green networking testbed; the ETOMIC high precision network measurement testbed; and the HEN emulation testbed. Associated with these testbeds were similarly successful control- and experimental-plane software. OpenLab advanced

these prototypes with key enhancements in the areas of mobility, wireless, monitoring, domain interconnections, and the integration of new technologies such as OpenFlow. These enhancements are transparent to existing users of each testbed, while opening up a diversity of new experiments that users can perform, from wired and wireless media distribution to distributed and autonomous management of new social interactions and localized services, going far beyond what can be tested on the current Internet. OpenLab's interoperability work brought FIRE closer to the goal of a unified Facility and provided models that were promoted to the Future Internet PPP. Finally, the project, through two open calls, supported users in industry and academia, notably those in FP7 Future Internet projects, who proposed innovative experiments using the OpenLab technologies and testbeds. Besides supporting users with a single portal and authentication mechanism, providing a direct access to their preferred testbeds, It opened an avenue for radically new needs covering the so-called verticals (smart cities, industrial Internet, transportation, environment, e-Health where several heterogeneous technologies have to be combined in a single experiment. This is a unique feature that is brought to the experimenters by OpenLab/OneLab.

12.2 Problem Statement

Experimentally-driven research is key to success in exploring the possible futures of the Internet. An open, general-purpose, shared experimental facility, both large-scale and sustainable, is essential for European industry and academia to innovate today and assess the performance of their solutions.

These were exciting times for those involved in creating new computing and communications applications, exploiting new technologies, and in seeing the world we live in change with the results. OpenLab aimed to play a key role in making these changes happen, and making Europe the hub for these changes.

Since computer applications now reach the home, the automobile, and the street, they go beyond making business and government services more efficient, and now form part of our social fabric. New ideas that start at the edge of the Internet, or of the telecommunications network, do not wait to be carefully deployed, or "rolled out" by industry, but are instead pulled out by users from App Stores, to be tried at modest or sometimes zero expense. However, their ultimate success or failure often depends on how well the infrastructure supports their requirements for interactivity and the

responsiveness with which the network as a whole delivers the data and rich media that the users expect and now require. To a greater extent than ever before, the usability and naturalness of an application controls its fate. That was a time in which the portfolio of testbeds developed through the Future Internet Research and Experimentation Initiative (FIRE) could have a major impact.

A second area in which the creation of new computer applications and the businesses that they support was changing rapidly was the Cloud. New services were developed and then deployed to businesses only as they were needed. The results were a promising source of growth for business of all sizes.

At the same time as we are seeing a wealth of services being deployed in data centres, we are seeing an ever-wider distribution of data generation and storage. Computing devices are getting smaller, giving impetus to an increasing use of data input by cameras, sensors, and crowd-sourcing from cell phones and vehicles. Local data storage is cheap, whereas communications are expensive and still slow, so we expect an increasing fraction of the world's information to remain close to where it is first captured and reduced to analyzable form. However, once made analyzable, these local, managed pools of data can be accessed and used from around the globe, wherever the ultimate users are found.

Innovations today come from all parts of the world. Skype is one example with its origins in Europe. Scandinavian firms have led the definitions of 3G and LTE, or 4G, mobile technologies. CERN has contributed the Grid approach to high performance, cost-effective computing. The German automobile companies play a leading role in automating personal transportation, and Europe is a world leader in rapid intercity trains. The USB memory stick was first productized in Israel. And in Japan last year, half of the best-selling popular novels were composed on smart phones or Blackberries during commute time. Note that these innovations tended to combine novel elements on more than one scale – for example, centralized information and smart phone applications with localized inputs, or more widely distributed data integrated by applications smart enough to exploit available resources that change as the user's location changes. But all these new ideas need better infrastructure to support them, need tuning based on understanding the properties of this new internet. FIRE's purpose is to provide the environment in which we can make the Future Internet happen on a small scale today, allowing innovators to tune while exploring these new ideas. OpenLab therefore put a special emphasis on linking testbeds of different types to prototype these fast growing applications.

12.3 Background and State of the Art

In this section, we provide a state-of-the-art for the topics addressed in OpenLab, at the time of the beginning of the project.

12.3.1 Federation in the Control and the Experimental Plane

Being able to stitch together several, potentially very different testbeds was at that time an area of very active research and development. The SFA architecture and OMF were two approaches for federation developed in the OneLab project. GENI had been pursuing this goal as well, advancing a set of four **control frameworks**, each with its own approach to federation [1]. Two of these clusters, namely PlanetLab (cluster B) and ProtoGeni (cluster C) were based on the SFA architecture. Although their initial implementations were not interoperable, an ongoing effort was carried out in this direction. Built on top of a secure layer for issuing API calls, SFA defined basically three kinds of services, namely a Registry that provides an index to the resources known to the system, an Aggregate Manager that manages the target testbed, and a Slice Manager that performs routing and aggregation of data in the meshed federation. SFA is highly decentralized, in that ideally a given user, although registered at any one authority in the federation, is able to browse and allocate resources from the entire mesh. In order to be able to cope with any sort of testbed, including the ones that have not yet been designed or invented, SFA makes no assumption as to the actual resource description languages that are left to the underlying testbeds and simply forwarded through the control plane. The ORBIT (cluster E) paradigm had a rather different approach, as security concerns were traditionally less crucial in this environment, where access control can be safely implemented at a single entry point for each testbed. OMF, the software behind ORBIT and many more testbeds worldwide, had been embracing a more holistic approach which combines the control, experimental, and measurement plane through a common set of design principles underpinning its suite of tools and services. One distinguishing feature in the context of the control plane is OMF's focus on efficient use of resources, as wireless testbeds tend to be owned by single entities that have strong economic incentives to have their users be as efficient as possible when running experiments. As an outcome of the Panlab/PII project, Teagle [2] built upon a resource federation framework to control distributed, highly heterogeneous resources. The model was more centralized compared to SFA, in the sense that a common information model allowed the detailed description of resources. Centralized services such as the orchestration and provisioning of federated

resources make extensive use of the common resource descriptions to allow for different granularities of resource abstraction. The aim was to support the user in working with the federated resources, exposing the necessary level of detail without overwhelming him with complex configuration requirements. Teagle as a trusted entity controls resources like physical and virtual machines, devices, and software across pan-European testbeds.

Concerning **experiment control**, it turns out that various user tools could exist that supported methodologies of defining experiments and best practices to conduct experiments. Such description of an experiment contained: i) resource requirements: what kind of resources are needed for the experiment, ii) resource configuration: user defined parameters applied on a resource, iii) resource relationships: resources might publish to or consume data from other resources, iv) workflow information: describe the needed provisioning tasks to be performed in order to create the experiment.

Regarding data storage and federation of data, the nmVO2 [3] represented a significant leap ahead in the state of the art infrastructures (DatCat, PERSONAR, MOME, etc.) that was storing meta-information of the monitoring and measurement data and was returning a pointer to a zipped file hosted by the owner of the data.

12.3.2 Wireless Testbeds

At the time of OpenLab many wireless testbeds for evaluating algorithms and protocols and validating communication techniques had been deployed. The most widely known were ORBIT [4] (WiMAX, Wi-Fi, Bluetooth, Zigbee and cognitive radio), MIT's Roofnet [5] and WARP [6] at Rice University, which focused on software-defined radio. These testbeds were open to the research community; however, they were limited by design drawbacks (such as the focus on static configurations, the very diverse resource descriptions, the very specific use policies) that prevented users from exploiting the many interesting features in an efficient way and that hinders tested facility owners to provision their facilities with better utilization of their resources. Other testbeds such as DieselNet at UMass [7] and the EU N4C testbeds [8] were focusing on more disruptive technologies such as delay-tolerant and opportunistic networking. Those testbeds were however closed to external experimenters. The OneLab community aimed to extend those testbed's initiatives by supporting a better framework for management and scheduling. NICTA [9], WINLAB [10], UTH [11] and other institutions had collaborated to develop and adopt a more efficient scheme for testbed management and control, based on their needs.

As a result OMF, a framework for unified testbed management, and the NITOS scheduler, a resource reservation scheduler providing slicing features, were developed, giving the opportunity to wireless network researchers to experience a more efficient and user-friendly experimental environment. Several institutions worldwide had adopted OMF for their testbeds.

12.3.3 Wired and Emulation Testbeds

Several EU funded projects promoted OpenFlow at the beginning of OpenLab, like OFELIA [12] which created an experimental facility based on OpenFlow. Five interconnected islands based on OpenFlow infrastructure had been created to allow experimentation on multi-layer and multi-technology networks. EU FP7 CHANGE explored the capabilities of OpenFlow to develop architecture for flow processing platforms within the network and individual processing of different flows. Additionally it tried to develop on-path and off-path flow processing and be the basis for flexible deployment of innovative services. The FP7 project SPARC [13] aimed at implementing a new split in the architecture of Internet components. The project would investigate splitting the traditionally monolithic router/switch architecture into separate forwarding and control elements. This functional split supports network design and operation in large-scale networks with multi-million customers who require a high degree of automation and reliability.

In the domain of the OpenFlow technology, there was clear orientation towards modelling the OpenFlow protocol functionality so that it could be offered to the wired platform users as a collection of network resources. As the wired platform would be a collection of physical and virtual network infrastructures, OpenFlow resources might correspond to a physical network infrastructure, a virtual one or a mixture of both.

Another important field addressed with new testbed enhancements is the domain of media streaming applications. As of today, experimentation on media streaming has been restricted in the field of research for coding algorithms with the scope to lighten the traffic volume of multimedia content. The VITAL++ FIRE project¹ has demonstrated the immense interest of the wider ICT community for experimentation on P2P multimedia content routing with optimum use of network resources. In the frame of the VITAL++ project, research has focused on demonstrating the feasibility of accommodating some particular P2P routing algorithms in an IMS testbed. The main challenge set in this project was the design of a generic mechanism for P2P algorithms incubation across testbed networks.

12.4 Approach

OpenLab's proposition was to bring together the essential ingredients for an open, general-purpose, shared experimental facility, both large-scale and sustainable. OpenLab objective was to build and open the OneLab facility. We wanted to extend early prototypes of testbeds, middleware, and measurement tools so as to provide more efficient and flexible support for a diverse set of experimental applications and protocols. The prototypes include a set of demonstrably successful testbeds: PlanetLab Europe, with its 153 partner/user institutions across Europe; the NITOS and wiLab.t wireless testbeds; two IMS (IP Multimedia Subsystem) telco testbeds for exploring merged media distribution; a green networking testbed; the ETOMIC high precision network measurement testbed; and the HEN emulation testbed. Associated with these testbeds were similarly successful control- and experimental-plane software. OpenLab wished to advance these prototypes with key enhancements in the areas of mobility, wireless, monitoring, domain interconnections, and the integration of new technologies such as OpenFlow. These enhancements were planned to be transparent to existing users of each testbed, while opening up a diversity of new experiments that users could perform, extending from wired and wireless media distribution to distributed and autonomous management of new social interactions and localized services, going far beyond what could be tested on the current Internet. OpenLab results will advance the goal of a unified Future Internet Research and Experimentation (FIRE) facility. Finally, OpenLab issued open calls to users in industry and academia to submit proposals for innovative experiments using the OpenLab's technologies and testbeds, and devoted one million euros to funding the best of these proposals.

OneLab came from a vision originated in 2005, built on several issues related to experimentally driven research. The networking community was facing a few successes in its ability to build testing tools (like PlanetLab or Emulab) but many more failures due to well-identified causes. In addition, a challenge that is still open for our community is to develop reproducible research, meaning that one should be able to reproduce the results that are published and supports a discovery.

This vision considered that an experimenter, namely, the one that was using the facility, should had access to an ecosystem or a "market" of various resources managed by different authorities. For this purpose, the experimenter would register to one such authority that would act as a mediator towards its peers.

The beauty of this model was grounded on the observation that there existed plenty of valuable resources out there that one could benefit if an open access was provided. Some of these resources might be unique, or the sum, or combination of them might be valuable. As it might not be the role of the resource owner to manage the users, this was delegated to an authority according to some constraints and obligations. In addition, it became quickly evident there is not a single testbed that fit all needs and that, solely, a federated model would succeed to embrace the vision (Figure 12.1).

OneLab project was one of the pillar of the European FIRE1 initiative and the initiator of the federation concept.

Enabling this vision required to define an architecture that supported the underlying concept of federation that was originally introduced in OneLab. Federation empowers to run services and tests using resources provided by autonomous organizations. Three main technology accelerators were identified:

- Virtualization,
- Open Source,
- Open Data.

Virtualization allows synthetic polymorphism (diversity of technologies) from one platform. In addition, it can create policy and security boundaries

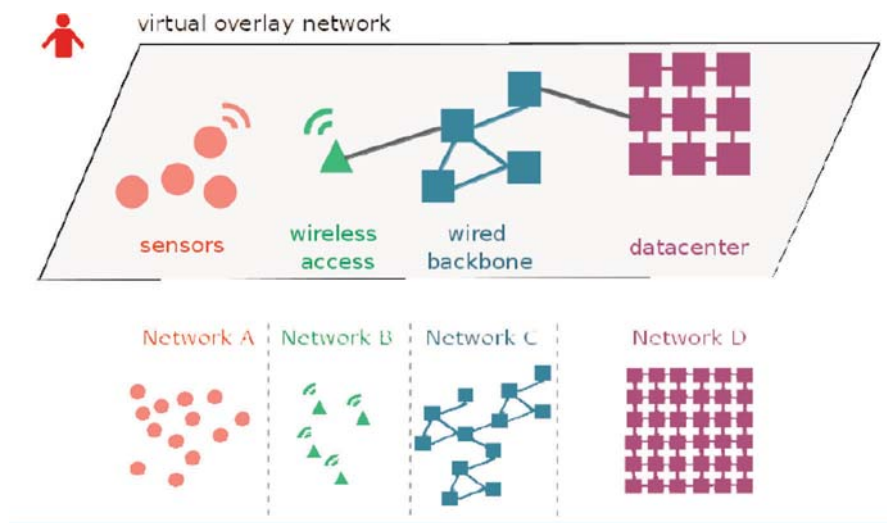


Figure 12.1 The federation of heterogeneous resources (provided by testbeds).

that are not the same as physical boundaries. **Open Source** was criticized at the beginning of OneLab as it did not support a given business model to pay for the usage of resources. Nevertheless, this approach has been proved credible in order to build a community of contributors to support the architecture. This does not even preclude any business model for the future use of the facilities. Finally, **Open Data** has not been given high-priority but will become dominant in the future as the first class citizen in any experiment is the data that have been produced and our ability use, document, and eventually share them.

Therefore, it became instrumental to address the following questions:

- What is the right level of abstraction, the minimum set of functionalities to be adopted to share resources owned by various authorities?
- What is the governance model that best supports subsidiarity?
- And finally, is there a business model or how can we enforce sustainability?

OpenLab’s work in all these areas was assessed at two interoperability testing events (also known as “bakeoffs” or “plugfests”), that were planned to be the occasions to see the extent to which the tools worked across the testbeds.

12.5 OpenLab Prototypes

This section describes the prototypes, a set of existing tools and testbeds, as of January 2011, illustrating the maturity of many of these technologies and their readiness to be integrated into a larger facility.

The prototypes described here emerged mostly from earlier FIRE efforts, notably the OneLab2 and PII projects and the work of IBBT. In addition, OpenLab brought in a number of excellent contributions from elsewhere, such as UCL’s HEN testbed, which was developed with UK national funding.

We organised our work on tools into two categories: control plane tools and experimental plane tools. Control plane tools largely work behind the scenes to support basic testbed operations and federation, whereas experimental plane tools were visible to the user and depend upon the control plane in order to function. The distinction between these two planes was similar to the notions of kernel and user space in the operating systems arena.

The tools that OpenLab started with were each typically specific to a single testbed environment. Our embrace of multiple tools with similar capacities was intentional: each had a particular way of doing things that would be attractive

to a particular type of user with particular needs. Indeed, some tools such as the OMF [9] Experiment Controller and PLE's MySlice [14] interface came with existing communities of hundreds of users who were already comfortable working with them. OpenLab extended these tools' coverage. By requiring a heterogeneous set of tools to function across multiple testbeds we put our emerging interoperability standards to the test.

Other prototypes were individual testbeds, or were tools that were specific to a testbed or type of testbed. We grouped these prototypes into two broad categories, wireless and wired, corresponding to two broad research networking communities. The groupings were not strict: some technologies described in one category could also be applied in another.

12.5.1 Wireless Prototypes

12.5.1.1 NITOS (Network Implementation Testbed using Open Source code)

NITOS [15] is an OMF-based wireless testbed in a campus building at UTH in Volos, Greece. It consisted of 45 nodes equipped with a mixture of Wi-Fi and GNU-radios, as well as cameras and temperature and humidity sensors. Two programmable robots provided mobility. This publicly available testbed supported experiments across all networking layers. In addition to OMF, the testbed employed locally developed tools: the NITOS scheduler, a resource reservation application, and TLQAP, a topology and connectivity monitoring tool.

12.5.1.2 w-iLab.t

The w-iLab.t testbed [16] is a wireless mesh and sensor network infrastructure deployed across three floors of the IBBT office building in Ghent, Belgium. It contained 200 locations, each equipped to receive multiple wireless sensor nodes and two IEEE 802.11a/b/g WLAN interfaces. Wi-Fi and sensor networks operated simultaneously, allowing complex and realistic experiments with heterogeneous nodes and multiple wireless technologies. In addition, shielded boxes used to accommodate nodes that can be connected over coax cables to RF splitters, RF combiners and computer controlled variable attenuators, thus allowing fully reproducible wireless experiments with emulated dynamically changing propagation scenarios. With an in-house designed hardware control device, unique features of the testbed included the triggering of repeatable digital or analogue I/O events at the sensor nodes, real-time monitoring of the power consumption, and battery capacity emulation.

12.5.1.3 DOTSEL

The DOTSEL testbed at ETH Zürich is focused on delay-tolerant opportunistic protocols and applications. It was composed of 15 Wi-Fi equipped Android Nexus One devices that were carried by staff members, and five Wi-Fi a/b/g ad-hoc gateways.

12.5.2 Wired Prototypes

12.5.2.1 PLE (PlanetLab Europe)

PLE [17] is the European arm of the global PlanetLab system, the world's largest research networking testbed, which gives users access to Internet-connected Linux virtual machines on over 1000 networked servers located in the United States, Europe, Asia, and elsewhere. Nearly 1000 scientific articles mention the PlanetLab system each year, including papers in such prestigious networking and distributed systems conferences as ACM SIGCOMM, ACM CoNEXT, IEEE INFOCOM, ACM HotNets, USENIX/ACM NSDI, ACM SIGMETRICS, and ACM SIGCOMM IMC. Researchers use PLE for experiments on overlays, distributed systems, peer-to-peer systems, content distribution networks, network security, and network measurements, among many other topics.

Established in 2006 and developed by the OneLab initiative, PLE is today overseen by four OpenLab partners: UPMC, INRIA, HUJI, and UniPi. UPMC handles testbed operations and INRIA co-leads, along with Princeton University, development of MyPLC, the free, open-source software that powers PlanetLab. The PlanetLab Europe Consortium has 150 signed member institutions: mostly universities and industrial research laboratories, each of which hosts two servers that it makes available to the global system. These institutions are home to 937 users. On a typical recent day, 244 were connected to on-going experiments.

OpenLab extends both the PlanetLab software and the PlanetLab Europe Consortium.

12.5.2.2 HEN (Heterogeneous Experimental Network)

HEN [18], built between 2005 and 2010 by UCL, provides 100 server-class machines with between 6 and 14 NICs each, interconnected by a Force10 E1200 switch with 550 Gigabit ports and 24 10-Gigabit ports. This infrastructure allows emulation of rich topologies in a controlled fashion over switched VLANs that connect multiple virtual machines running on each host. The precise control of topology and choice of end-host operating

system possible on HEN are particularly valuable facilities to networking and distributed systems researchers.

Many dozens of researchers actively use HEN: at Stanford University, the University of Lancaster, NYU, the Nokia Research Centre, and NEC Labs Europe, to name a few. UK- and EU-funded projects, including the EPSRC-funded Virtual Routers project, EPSRC-funded ESLEA project, EU FP7-funded Trilogy project, and EU FP7-funded CHANGE project, have all generated the bulk of their experimental results on HEN. Results have been published in prestigious networking and distributed system venues including ACM SIGCOMM, ACM HotNets, USENIX/ACM NSDI, USENIX Security, ACM CCR, ACM CoNEXT, Presto, FDNA, PMECT, ICDCSW, and LSAD.

12.5.2.3 The WIT IMS testbed

The TSSG/WIT NGN IMS testbed [19] is an Irish nationally-funded initiative serving telecom firms seeking to develop or test NGN services. It provides them with advanced multimedia services, such as conference calling and handling of presence information. The testbed is a carrier grade NGN platform based on the Ericsson IMS Communications System (ICS). The SIP based horizontal network architecture includes an Ericsson IMS core and the components for managing sessions, addressing, subscriptions and IMS interworking components with the relevant gateways for connectivity to other networks. The testbed has recently been upgraded with pico/femto cells to allow secure remote access to the test facility. The network also includes support systems for handling provisioning, charging, device configuration and operation and maintenance.

Clients include IP centrex companies, a location based service provider, and developers of pico/femto cell technology. International customers have conducted testing in the area of IMS security and testbed interconnection using the GSMA Pathfinder service operated by Neustar.

12.5.2.4 The University of Patras IMS testbed

The University of Patras IMS testbed supports PSTN testing scenarios: calls between a PSTN network and any PSTN number (including international and mobile numbers); and calls between IP phones (either soft phone or hard phone) and any PSTN number (including international and mobile numbers). The testbed has been used in numerous interoperability experiments with the carrier grade network of Telecom Austria, and the NGN testbeds of Siemens AG in Munich and Telefónica TID in Madrid. It is currently hosts experiments from the FP7 VITAL++ project. Integration of the testbed into the Teagle

framework was carried out under the PII project. In OpenLab, the testbed will be enhanced to incubate P2P/NGN QoS reservation algorithms and establish experimentation paths taking advantage of the OpenFlow protocol.

12.6 Technical Work

12.6.1 Federation in the Control and the Experimental Plane

Research in the networking area has fostered the emergence of a wide variety of experimental testbeds. The vision, that OpenLab had been promoting from its very beginning, had it that disruptive innovations in the networking area would emerge from giving researchers easy access to all these resources in an open and consistent way, thus creating opportunities to conjugate all the new capabilities at a large scale.

OpenLab's activities had been instrumental in bringing this vision to reality, by tackling the general issue of testbed federation. Our achievements in this field were very substantial, both at the design and implementation levels. On the design front, OpenLab had been an active contributor to the architecture and specification of SFA, that offers a common way for testbeds and tools to expose, discover and provision resources – what we call “Control Plane” – in an homogeneous way, and that was defined inside an international community that, despite being informal, has representatives from virtually any kind of networking testbed in the developed countries. Still on the design side, OpenLab had proposed **FRCP**, a testbed-neutral layer for managing live resources – what we call “Experimental Plane” – to serve the same kind of purpose as SFA but during experimentation and not only in the preparation phases. FRCP has likewise reached a very wide consensus over the community and is starting to be widely available.

Like always when general adoption is at stake, proposing specifications and architecture is not enough if it does not come with at least one reference implementation. This is why OpenLab developed **SfaWrap** [20] and **OMF6** [9] that provided such a reference implementation of SFA and FRCP respectively, see also Figures 12.2 and 12.3.

Building on top of this architectural foundation, OpenLab had created a legal framework for operating a wide and heterogeneous federation of testbeds that spanned beyond OpenLab per se, and that we had named **the OneLab Consortium** (more details in Section 12.7.2). Starting with the OneLab Portal (see Section 12.7.2.2), researchers can enjoy all the benefits of testbeds federation on for example PlanetLab Europe that operates old-school

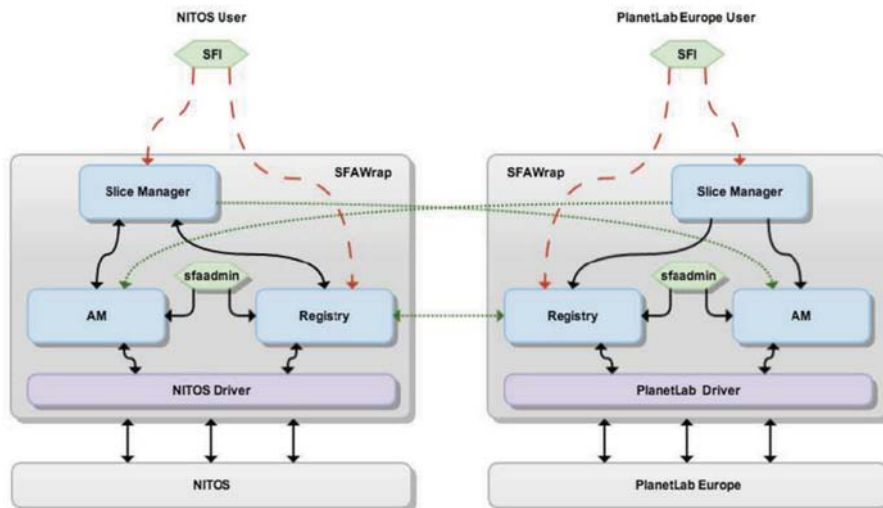


Figure 12.2 NITOS and PlanetLab Europe federation via SFAWrap.

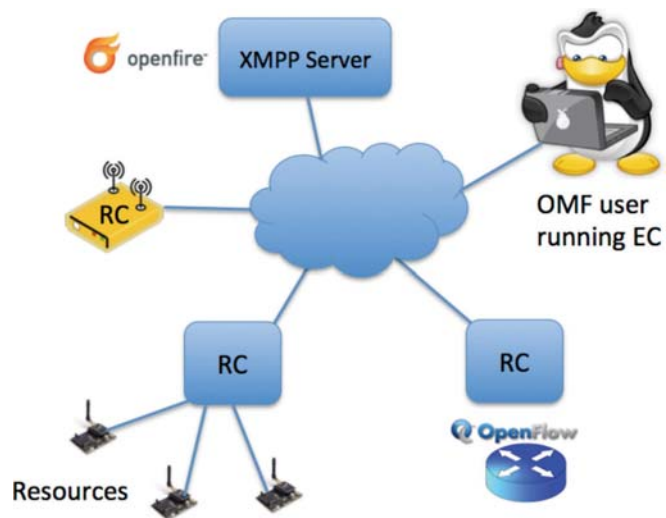


Figure 12.3 OMF 6 architecture.

wired servers, IoTLab that offers sensor nodes, NITOS that features WiFi nodes; several other European testbeds are in the process of joining.

This portal features higher level, more experimenter-oriented tools, that were developed within OpenLab; in this category let us quote **MySlice** [14],

alongside with its companion **manifold**, that runs under the hood in the OneLab Portal (see Section 12.7.2.2), and that offers a set of web-based tools for dealing with heterogeneity, in terms of both resources and measurements, but in a uniform manner, as depicted in Figure 12.4.

However, using the portal was not the only option, and third-party tools had also been implemented, that directly took advantage of SFA and FRCP to offer alternative all-in-one tools for researchers, like for example **NEPI** [21, 22].

To summarize, OpenLab has created a complete paradigm for deploying a federation of testbeds, and is now operating the OneLab Portal as a first production-grade such federation. We are hoping to provide valuable help to the research community thanks to this new tool, and are confident that the conceptual assets of OpenLab will be further enhanced by on-going and future projects, like Fed4Fire.




Figure 12.4 MySlice and manifold architecture.

Concerning experiment control, the most visible outcome is all the set of enhancements and added features brought to NEPI for supporting generically SFA and FRCP; these capabilities obviously extended the tool's scope drastically, as a very welcome addition to the already existing testbed-specific methods for accessing resources (like e.g. raw ssh). Concerning interoperability of measurement data, the MOI ontology was concluded with participation of TUB, iMinds and UAM. A data federation tool called GrayWulf for SQL data sources was developed that enables users to access various databases through a unified SQL-based querying interface. Both approaches integrated TopHat, EtomicDB and nmVO databases operated by different OpenLab partners. Concerning usage control, a detailed usage accounting activity was performed, for each of the testbeds participating in OpenLab.

Finally, two experiments were conducted through the two Open Calls of the project using and validating the activities of this section:

- **SNIFFER Experimentation:** a replicable base for long-running service using OpenLab and PlanetLab environment in order to better observe and track the long-term growth of various Storage Networks (Grids, Clouds, Content Delivery Networks, Information-Centric Networks) [23].
- **ECLECTIC Experimentation:** a new tool for testbed management for Peer-to-Peer (P2P) applications which includes improved support for resource allocation, deployment and state-of the-art monitoring over a range of experimental testbeds.

12.6.2 Wireless Testbeds

In the context of OpenLab, three existing FIRE wireless testbeds (NITOS, w-iLab.t and DOTSEL) were enhanced in terms of hardware with features like: LTE and WiMAX technologies, new wireless fixed and mobile nodes and directional antennas and in terms of software with the integration of the above hardware additions into the existing control and management tools. NITOS [24] facility was greatly extended with an indoor testbed featuring new powerful wireless nodes and directional antennas, a WiMAX testbed, an LTE testbed and an uncontrolled mobility testbed comprising of mobile phones carried by volunteers, as depicted in Figures 12.5 to 12.7.

The w-iLab.t [16] testbed was also extended with a number of wireless nodes and more importantly with a real life mobility testbed. In this testbed, wireless nodes were mounted on top of robots and are provided to the community for testing mobility issues, through user-friendly graphical interfaces and tools. The experimenter is able to use a web graphical interface, which enables



Figure 12.5 Extensions of the NITOS testbed: Icarus nodes on the left and directional antennas on the right.



Figure 12.6 WiMAX/LTE Base Station in NITOS testbed.

him/her to draw the desired path of the robots, simulate the scenario prior its execution or auto-detect collisions between the mobile nodes. Moreover, the testbed's administrator is able to monitor vital metrics like the exact status of each robot (docked, idle, active etc.), its remaining power and the access point that is connected, through the aforementioned GUI.



Figure 12.7 Mobile robots in w-ilab.t testbed.

Regarding the software developments of the wireless testbeds during the project, some of the main achievements are the extensions built for OMF, which is the framework that OpenLab and most of the wireless testbeds in the world use to orchestrate their testbed's resources. All the hardware extensions mentioned above, are now OMF compatible, by developing Resource Controllers (RC) for the WiMAX testbed, for the mobile phones of the uncontrolled mobility testbed and for the mobile nodes of the controlled mobility testbed enabling the user to control those resources through OMF. This way, the experimenter is able to design, develop and deploy an experiment using for example the WiMAX Base Station, some wireless nodes and some mobile phones, where he/she is able to test an Android application developed by him/her, or provided by the testbed. The most important aspect of this is that all these heterogeneous resources can be handled through one single OMF script, regardless of their physical location, as depicted in Figure 12.8.

Finally, the experimenter is able to design and deploy complex experiments on top of more than one wireless testbeds (for example NITOS and w-iLab.t) and take advantage of the diverse federation aspects enabled through the OMF framework. This way all the hardware extensions happened in the context of the project were integrated into OMF, providing an SFA interface, namely the necessary hook for the connection and the communication with other federation frameworks.



Figure 12.8 Demonstration of a complex experiment controlling heterogeneous resources with a single script.

Finally, two experiments were integrated in the context of this section's activities through the two Open Calls:

- SAVINE experiment: a Social-aware Virtual Network embedding framework for a wireless content delivery scenario [25].
- EXPRESS experiment: an innovative resilient SDN system able to withstand attacks, failures, mistakes, natural disasters and able to keep operating also in fragmented and intermittently connected networks [26].

12.6.3 Wired Testbeds

Regarding wired testbeds, OpenLab focused on the evolution of the four wired testbeds (PLE, OSIMS, WIT and HEN) participating in the project towards the support of a number of new features including: Software Defined Networking (SDN) capabilities, implementation of QoS support mechanisms, multi-homing functionality and finally support of testbed interconnectivity both from control and networking perspective to allow seamless usage of testbeds. SDN features were implemented, as planned, on the basis of support of OpenFlow. The OpenVSwitch was ported and installed in the testbeds and also integrated with each testbed's control framework and procedures for the provision of additional functionality on top of it (e.g policy enforcement,

resource discovery, and network overlay support). QoS support mechanisms were realised in terms of both P2P algorithms incubation and policy enforcement via network flow management mechanisms integrated in the overall control and management framework of the involved testbeds (OSIMS and WIT). Multi-homing was based on the enhancements applied to the HEN testbed with the addition of two external Internet links and the support of external access to the nodes by implementation of additional firewall and addressing features. Interconnectivity of testbeds proceeded along two axes. The first one related to the support of SFA mechanisms in testbeds where there was already a provisioning framework to be adapted so that resource discovery, reservation and configuration can be applied according to the SFA principles. The second one related to the networking interconnectivity for the support of experiments involving more than one testbed. The adopted approach was based on the use of Layer 2 Overlays that can be instantiated by use of the implemented OpenFlow enhancements. Control and management procedures as well as provisioning mechanisms in each testbed were updated and enhanced to support the network interconnectivity approach.

Finally, four experiments were integrated in the context of this section's activities through the two Open Calls:

- ALLEGRA: deployment and “proof-of-concept” testing of a lightweight greedy geographical routing algorithm.
- ANA4IoT: extension of the OpenLab/FIRE testbeds by mixing the available infrastructure (physical or virtual) to build a new scenario for testing different approaches and evaluate the capabilities to cope with IoT requirements.
- PSP-SEC: evaluation of a PSP SDN application running on top of OPENER, with the aim of delivering security in a way that BGP is not even aware that it is being secure.
- WONDER: experimentation with and evaluation of WebRTC service delivery mechanisms namely IMS and Web service delivery approaches.

The experiments have helped both experimenters and testbed owners to collect valuable feedback and focus better on issues relating either to their experimentation aspects or to their testbed mechanisms respectively.

12.7 Results and/or Achievements

A major achievement of the OpenLab project is the opening of the OneLab facility, the first open federation of heterogeneous testbeds, making OneLab sustainable and independent from the OpenLab project ending in September

2014. OneLab materialized the efforts of the OpenLab consortium during the course of the project. It also clearly stressed our ambition to deliver a Service as the main contribution of the project and not to focus on the underlying technology, although very important.

In section 12.7.1 we describe the particular outputs of the project and in section 12.7.2, we present the major outcome of the project, namely the OneLab Experimental Facility.

12.7.1 OpenLab Main Outputs

We benefited from the experience in architecting the Internet to design our architecture model. It is grounded on two principles:

- The “Hourglass” model of the Internet that identifies the IP protocol as the convergence layer. We’ll define one for the Federation of testbed resources;
- The peering model of the Internet that relies Customer and Providers and define peering agreements in a way that there is not a single point of control. Here, we will clearly identify Experimenters, Testbed owners or providers and the Facility itself that rule them all.

We therefore have defined the following abstractions:

- Resource: Testbed ensures proper management of nodes, links, switches.
- User/Experimenter: Testbed guarantees the identity of its users.
- Slice: A distributed container in which resources are shared (sharing with VMs, in time, frequency, within flowspace, etc.). It is also the base for accountability.
- Authority: An entity responsible for a subset of services (resources, users, slices, etc.).

SFA (Slice-based Federation Architecture) was designed as an international effort, originated by the NSF GENI framework, to provide a secure common API with the minimum possible functionality to enable a global testbed federation.

The fundamental components for testbed federation were built incrementally, as the understanding about the requirements became better understood. The first international realization of federation arose in 2007, as a mutual investment from PlanetLab Central, managed by Princeton, and Planet Europe, established by UPMC and INRIA in Europe. It was then, enlarged to both private and public instances of PlanetLab, allowing a user registered under one of these authorities to benefit from resources own by any other authority.

Nevertheless, resources were homogeneous and the usability was tight to PlanetLab users. It then became of utmost importance to enlarge and extend the federation principle to other type of resources, a more scalable model of federation and an increased ease of use. In parallel, started the important effort to complement and populate the architecture with components mandatory for the entire experiment life cycle.

The experiment lifecycle comprises the following steps:

1. User account & slice creation
2. Authentication
3. Resource discovery
4. Resource reservation & scheduling
5. Configuration/instrumentation
6. Execution
7. Repatriation of results
8. Resource release

Step 1 is handled by the Home Authority of the User, the one the user has registered with. Steps 2 to 4 and 8 concern all involved authorities. Steps 5 to 6 are not in SFA but other components such as OMF have been developed for this purpose. OMF is a control, measurement and management framework that was originally developed for the ORBIT wireless testbed at “Winlab, Rutgers University”. Since 2008, OMF has been extended and maintained by NICTA and UTH as an international effort.

SFA provides a secure API that allows authenticated and authorized users to browse all the available resources and allocate those required to perform a specific experiment, according to the agreed federation policies. Therefore, SFA is used to federate the heterogeneous resources belonging to different administrative domains (authorities) to be federated. This will allow experimenters registered with these authorities to combine all available resources of these testbeds and run advanced networking experiments, involving wired and wireless technologies.

Another component of the SFA layer is the Aggregate Manager (AM), which is required in each SFA-compliant testbed. The AM is responsible for exposing an interface that allows the experimenters to browse and reserve resources of a testbed. The SFA AM exports a slice interface that researchers interact with to set up, control, and tear down their slices. When the Control and Management Framework (CMF) of a testbed is not SFA-compliant, a so-called SFA driver is required to translate SFA originated queries into queries for the testbed. This driver wraps the CMF and exposes a standard interface

to the AM. SFA Wrap [20], was then designed to ease the adaption of SFA by testbed owners so that they only have to develop the part related to the specificities of their own testbeds. This is a shared development model that has been widely adopted by the testbed community.

The SFA layer is composed of the SFA Registry, the SFA AMs and drivers. The SFA Registry is responsible to store the users and their slices with the corresponding credentials.

MySlice [14] was introduced as a mean to provide a graphical user interface that allows users to authenticate, browse all the testbeds resources, and manage their slices. This work was important to provide a unified and simplified view of many hidden components to the experimenter. At the same time, it provides an open environment for the community to enrich the portal through various plugins specific to each testbed or environment. The basic configuration of MySlice consists on the creation of an admin user and a user to whom all MySlice users could delegate their credentials for accessing the testbed resources. In order to enable MySlice to interact with heterogeneous testbeds, MySlice has to be able to generate and parse different types of RSpecs (Resource Description of the testbeds); this task is performed by plugins. MySlice has been widely adopted by the community and is currently an international effort. As of today MySlice has been adopted by the following testbeds (or Projects): FIT (France), F-Lab (France), FanTaaStic (EU), Fed4Fire (EU), OpenLab (EU), FIBRE (Brazil), FORGE (EU), CENI (China), SmartFire (Korea) and III (Taiwan).

Finally, **Instrumentation and Measurement** has always been considered as a major component since the development of the Internet of testbeds concept. Indeed, core activities involved in experimentation are related to identifying, assessing and providing a set of tools and methodologies to create an empirical evidence base by measuring and adequately representing Internet data (Metrology) and, in a subsequent process, information (Mediametry) thanks to experimental investigation. The role of Instrumentation and Measurement is therefore strategic as it aims at providing the experimental validation and assessment of the scientific principles supported in the experiment. It is also involved in measuring unknown quantities, ranging from low-level parameters, such as packet loss, to high-level as individual user utilities and concerns. Today, every testbed has integrated more or less mature tools to support the user's experiment. Federating heterogeneous testbeds creates the necessity to organize measurement tools and methodologies, as well as consider a relevant architecture for this purpose. An important effort has also been dedicated by the community but has taken more time to materialize

although many tools are already available such as the Manifold framework in operation in the OneLab facility.

12.7.2 The OneLab Experimental Facility

A major achievement in OpenLab exploitation is the birth of a sustainable OneLab Experimental Facility [27], with its defined governance and management structure independent and beyond any projects' lifetime. The Consortium Agreement for the OneLab Experimental facility was signed in March 2014 by five OpenLab partners: UPMC, UTH, INRIA, iMinds and TUB. The OneLab Experimental Facility Parties are partners both from the OpenLab project and from the FIT project.

- Among the Parties are joint participants in the OpenLab project, which is funded under the Seventh EU Framework Programme for Research and Technological Development, and which includes as part of its work programme the establishment of a consortium to facilitate the use of testbeds for networked computer communications;
- Among the Parties are joint participants in the FIT project, which is funded under the French national Equipements d'Excellence programme, and which is currently building a group of testbeds for networked computer communications, and which has opted to promote the use of the FIT testbeds in the context of a larger consortium that includes other testbeds;
- The Parties having collaborated in these and other projects, and having also worked severally, in recent years on testbeds for networked computer communications, this work now having reached a state of maturity in which the testbeds can be presented collectively, or in a "federation", to users.

The Parties have agreed to create a consortium called OneLab, that, for the first time in this domain, facilitates the use of multiple testbeds for networked computer communications, exists independently of any individual project, and exists beyond the lifetime of individual projects;

1. Through OneLab, to create and manage the OneLab Experimental Facility, an overarching management structure with the technical systems necessary to support a group of testbeds, referred to as Affiliated Testbeds, which will have signed agreements with OneLab. This support is to be guided by a principle of "subsidiarity", whereby only those functions that are best performed by OneLab are allocated to the Consortium, all other

functions being reserved to the legal entities that manage each Affiliated Testbed;

2. Through OneLab, to facilitate the use of the Affiliated Testbeds by individuals who are connected to legal entities that have entered into signed Membership Agreements with OneLab, the legal entities thereby becoming Members. These individuals include, but are not limited to, researchers who are salaried by a Member, students who are enrolled at a Member that is an institution of higher education, and others who have a visitor's status with a Member; The Parties wish by means of this Agreement to set the terms and conditions for implementing the OneLab Consortium,

The OneLab Experimental Facility materialises in practical terms in an established governance structure of the OneLab Consortium and in OneLab Portal.

12.7.2.1 OneLab Consortium

OneLab Consortium Agreement defines the structure, the functions, the rights and responsibilities of Parties for governing and managing OneLab Experimental facility. The Consortium defines three bodies, which are: the Governing Board, which is the sole decision-making body of the Consortium; the Board of Affiliates, an advisory body that deliberates on issues relating to the Affiliated Testbeds; the General Assembly, an advisory body that deliberates regarding OneLab's plan of activities. These are, collectively, the Consortium Bodies. In the context of these bodies, three categories of participants in the Consortium are defined:

- Governors, which are the Parties to this Agreement;
- Affiliates, which are, principally, those legal entities responsible for Affiliated Testbeds;
- Members, which are, principally, those legal entities that bring users to the Affiliated Testbeds.

The role of the Coordinator is also defined, the coordinator being the legal entity that represents and manages OneLab. Furthermore, the roles of two officers appointed by the Coordinator are defined:

- the President, who carries out the tasks related to representing OneLab;
- the Executive Director, who carries out the tasks related to managing OneLab.

OneLab Experimental Facility welcomes Affiliates and Members to benefit from and to enlarge even more the impact of shared, federated, networked computer communication resources.

12.7.2.2 OneLab Portal

OneLab value proposition is the ease of use and support in experimentation. OneLab Portal [28] offers easy access to testbeds; through OneLab, the experimenter, (a researcher, a developer, and innovator) can easily test the software system in any one, or any combination of the following networked communication environments:

- Internet-overlaid testbeds
- Wireless, sensing, and mobility testbeds
- Broadband access and core testbeds
- Network emulation environments

OneLab aims to attract users beyond its own testbed providers and its immediate stakeholders, to offer resources also for e.g. educators, learners (FIRE FP7 project FORGE, EIT ICT Labs Master School) and SMEs. There is much to understand in networked communication testbeds: each platform's hardware capabilities, how the available software environments be configured and loaded onto a platform, the many features of the experiment control tools, etc. At OneLab, we offer a skilled team that is happy to assist throughout the experimentation cycle.

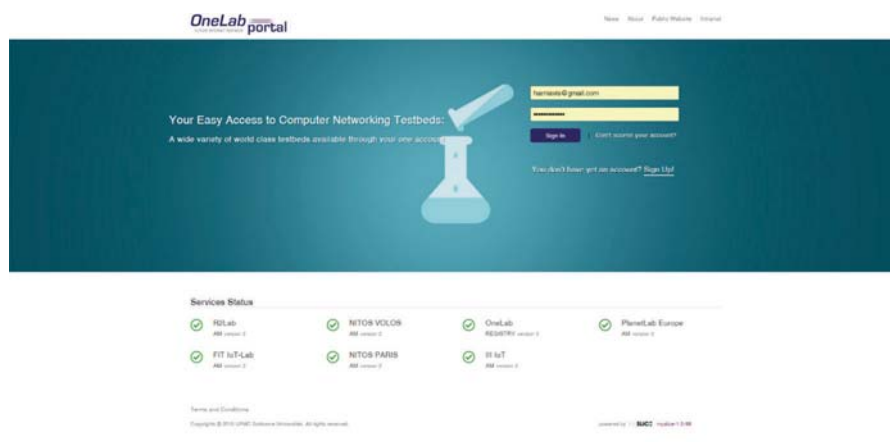


Figure 12.9 The OneLab Portal.

- For all users – academic and industrial: OneLab provides with online tutorials, documentation, and invitations to hands-on workshops and community events at which experimenters provide feedback and platform and tool developer describe their plans.
- For SMEs & other industrial users: OneLab team can accompany through the entire process of designing and running your experiment and interpreting its results.

OneLab provides tools that make it easy to use its testing environments:

- Through OneLab Portal one can access any of the testing resources. If the testing is repeated in more than one environment, there is no new account to open, no new system to learn.

Through OneLab Portal one can select highly capable experiment control tools. These are free open-source tools which are evolving to meet the needs of an ever-growing community of experimenters, and can be tailored, if needed, for particular requirements.

12.8 Conclusions

OpenLab has made a major contribution by deepening the capabilities of its various testbeds, inherited from FIRE's former OneLab and Panlab initiatives as well as other valuable sources. OpenLab advanced early FIRE prototypes that had proven their worth; this integrated infrastructure project associated and extended them, enhancing the value of the FIRE portfolio of facilities.

Combining these advanced infrastructures provided many examples of opportunity, with media distribution and localized services delivered to wireless clients perhaps capturing the most attention. As a result, our second major effort in OpenLab was to provide interoperability of these different services, allowing access and authorization to one to permit access to others, within the policies for usage and security required by each. Such common access methods form a control plane for the FIRE testbeds.

OpenLab wanted to strengthen the current offering by constructing a standard set of experiment deployment procedures, or a federated experimental plane, through which resources could be described, found, and reserved or allocated immediately. This also required implementing standards for the description of experimental configurations, for a real or simulated workload, and for the forms in which the resulting data will be logged, aggregated, analysed, and archived. Monitoring tools were critical for understanding usability issues that affect new applications and would also give us a means

of seeing how testbed utilization develops over time. For testbeds to be sustainable, they should evolve to meet new interests. Only sustained access to testbeds will permit the creation of long-running experimental services so that we can understand their strengths, their weaknesses and the degree to which they meet users' expectations. This sustainability has been one of the strengths of PlanetLab and of PlanetLab Europe in recent years.

The generic control and experimental planes introduced above also needed to be instantiated and extended in each of the wireless and wired testbeds included in this project. The chief tool we employed for wireless is OMF (cOntrol and Management Framework), a framework that is used to control around 20 testbeds worldwide, including several in Europe. The tools have been expanded to support both controlled and uncontrolled experimental conditions. The wired testbeds in OpenLab allow innovation within both the telco and data network paradigms (IMS and IP protocols). The availability of OpenFlow protocols further enriches the mix of activities that can be supported and the depth into the networking stack to which experiments can probe or prototype. The Heterogeneous Experimental Network (HEN) in operation at UCL brings the two types of environments closer together. SFA (Slice-based Federation Architecture) was deployed and extended as the envelope to federate the various technology-specific control and experiment planes.

OpenLab was instrumental to push the envelope of knowledge and tools in FIRE such that the OneLab Facility was successfully launched as an independent facility on August 2015, supporting a broad and diverse set of experiments.

References

- [1] "GENI Spiral 2," [Online]. Available: groups.geni.net/geni/wiki/SpiralTwo
- [2] "Teagle Portal," [Online]. Available: <http://www.fire-teagle.org/>
- [3] "nmVO," [Online]. Available: <http://nm.vo.elte.hu/>
- [4] "ORBIT Testbed Portal," [Online]. Available: <http://www.orbit-lab.org>
- [5] "MIT Roofnet," [Online]. Available: <https://en.wikipedia.org/wiki/Roofnet>
- [6] "WARP: Wireless Open Access Research Platform," [Online]. Available: <http://warp.rice.edu/>
- [7] "UMass DieaselNet," [Online]. Available: <https://dome.cs.umass.edu/umassdieselnet>

- [8] “N4C: Networking for Communications Challenged Communities,” [Online]. Available: <http://www.n4c.eu/>
- [9] “cOntrol and Management Framework,” [Online]. Available: <https://mytestbed.net/>
- [10] “WINLAB,” [Online]. Available: <http://www.winlab.rutgers.edu/>
- [11] “NITlab: Network Implementation Testbed Laboratory,” [Online]. Available: <http://www.winlab.rutgers.edu/>
- [12] “OFELIA experimental facilities,” [Online]. Available: <http://www.fp7-ofelia.eu/>
- [13] “SPARC FP7 project,” [Online]. Available: <http://www.fp7-sparc.eu/>
- [14] “MySlice,” [Online]. Available: <https://www.myslice.info/>
- [15] “NITOS Future Internet Facility,” [Online]. Available: <http://nitlab.inf.uth.gr/NITlab/index.php/nitos.html>
- [16] “w-iLab.t testbed,” [Online]. Available: <http://wilab2.ilabt.iminds.be/>
- [17] “PlanetLab Europe,” [Online]. Available: www.planet-lab.eu
- [18] “HEN, Heterogeneous Experimental Network,” [Online]. Available: mediatools.cs.ucl.ac.uk/nets/hen
- [19] “TSSG/WIT NGN IMS testbed,” [Online]. Available: ngntestcentre.com
- [20] “SFA Wrap,” [Online]. Available: <http://www.sfawrap.info>
- [21] A. Quereilhac, M. Lacage, C. Freire, T. Turletti and W. Dabbous, “NEPI: An integration framework for network experimentation,” in *19th International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, 2011.
- [22] “NEPI,” [Online]. Available: <http://nepi.inria.fr/>
- [23] A. Bak, P. Gajowniczek, M. Pilarski and M. Borkowski, “Automated Discovery of Worldwide Content Servers Infrastructure – the SNIFFER Project,” in *FedCSIS*, Warsaw, Poland, 2014.
- [24] “NITOS Future Internet Facility,” [Online]. Available: <http://nitlab.inf.uth.gr/NITlab/index.php/nitos.html>
- [25] A. Leivadreas, C. Papagianni and S. Papavasileiou, “Demonstration of a social aware wireless content delivery paradigm,” in *IEEE Infocom 2014*, Toronto, Canada, 2014.
- [26] S. Fdida, T. Korakis, H. Niavis, S. Salsano and G. Siracusano, “The EXPRESS SDN Experiment in the OpenLab Large Scale Shared Experimental Facility,” in *1st International Science and Technology Conference on Modern Networking Technologies: SDN & NFV*, Moscow, Russia, 2014.
- [27] “OneLab,” [Online]. Available: <https://onelab.eu/>
- [28] “OneLab Portal,” [Online]. Available: <http://portal.onelab.eu>

